

Partial Functional Correspondence

E. Rodolà¹, L. Cosmo², M. M. Bronstein³, A. Torsello², D. Cremers¹

¹TU Munich, Germany ²University of Venice, Italy ³University of Lugano, Switzerland

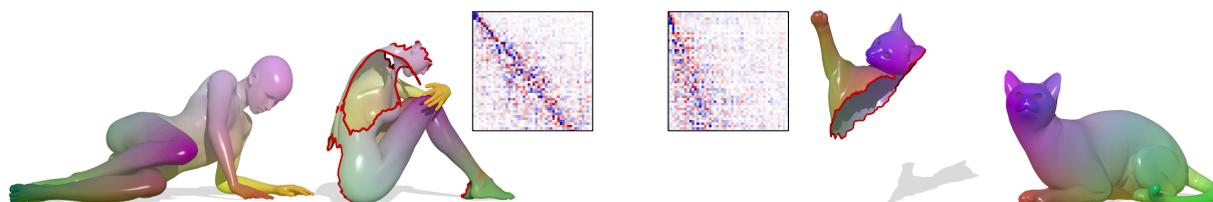


Figure 1: Partial functional correspondence between two pairs of shapes with large missing parts. For each pair we show the matrix \mathbf{C} representing the functional map in the spectral domain, and the action of the map by transferring colors from one shape to the other. The special slanted-diagonal structure of \mathbf{C} induced by the partiality transformation is first estimated from spectral properties of the two shapes, and then exploited to drive the matching process.

Abstract

In this paper, we propose a method for computing partial functional correspondence between non-rigid shapes. We use perturbation analysis to show how removal of shape parts changes the Laplace-Beltrami eigenfunctions, and exploit it as a prior on the spectral representation of the correspondence. Corresponding parts are optimization variables in our problem and are used to weight the functional correspondence; we are looking for the largest and most regular (in the Mumford-Shah sense) parts that minimize correspondence distortion. We show that our approach can cope with very challenging correspondence settings.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Shape Analysis

1. Introduction

The problem of shape correspondence is one of the most fundamental problems in computer graphics and geometry processing, with a plethora of applications ranging from texture mapping to animation [BBK06, KLCF10, KLF11, VKZHCO11]. A particularly challenging setting is that of *non-rigid correspondence*, where the shapes in question are allowed to undergo deformations, which are typically assumed to be approximately isometric (such a model appears to be good for, *e.g.*, human body poses). Even more challenging is *partial correspondence*, where one is shown only a subset of the shape and has to match it to a deformed full version thereof. Partial correspondence problems arise in numerous applications that involve real data acquisition by 3D sensors, which inevitably lead to missing parts due to occlusions or partial view.

Related work. For rigid partial correspondence problems, arising *e.g.*, in 3D scan completion applications, many versions of regularized iterative closest point (ICP) approaches exist, see for example [AMCO08, ART15]. Attempts to extend these ideas to the non-rigid case in the form of non-rigid or piece-wise rigid ICP have been explored in recent years [LSP08]. By nature of the ICP algorithm, these methods rely on the assumption that the given shapes can be placed in approximate rigid alignment to initiate the matching process. As a result, they tend to work well under small deformations (*e.g.*, when matching neighboring frames of a sequence), but performance deteriorates quickly when this assumption does not hold.

For the non-rigid setting, several metric approaches centered around the notion of minimum distortion correspondence [BBK06] have been proposed. Bronstein *et al.* [BB08,

[BBBK09] combine metric distortion minimization with optimization over matching parts, showing an algorithm that simultaneously seeks for a correspondence and maximizes the *regularity* of corresponding parts in the given shapes. Rodolà et al. [RBA*12] subsequently relaxed the regularity requirement by allowing sparse correspondences, and later introduced a mechanism to explicitly control the degree of sparsity of the solution [RTH*13]. Finally, in [SY14] the authors proposed a voting-based formulation to match shape extremities, which are assumed to be preserved by the partiality transformation. Being based on spectral features and metric preservation, the accuracy of the aforementioned methods suffers at high levels of partiality, where the computation of these quantities becomes unreliable due to boundary effects and meshing artifacts. Furthermore, these methods suffer from high computational complexity and generally provide only a *sparse correspondence*.

Pokrass et al. [PBB13] proposed a descriptor-based partial matching approach where the optimization over parts is done to maximize the matching of bags of local descriptors. The main drawback of this approach is that it only finds similar parts, without providing a correspondence between them. Windheuser et al. [WSSC11] formulated the shape matching problem as one of seeking minimal surfaces in the product space of two given shapes; the formulation notably allows for a linear programming discretization and provides guaranteed continuous and orientation-preserving solutions. The method was shown to work well with partial shapes, but requires watertight surfaces as input (e.g., via hole filling). Brunton et al. [BWW*14] used alignment of tangent spaces for partial correspondence. In their method, a sparse set of correspondences is first computed by matching feature descriptors; the matches are then propagated in an isometric fashion so as to cover the largest possible regions on the two shapes. Since the quality of the final solution directly depends on the initial matches, the method is understood as a “densification” method to complement other sparse approaches. Other recent works include the design of robust descriptors for partial matching [vKZH13]. In the context of collections of shapes, partial correspondence has been considered in [VKTS*11, CGH14, CRA*16].

All the aforementioned works are based on the notion of point-wise correspondence between shapes. Recently, Ovsjanikov et al. [OBCS*12] proposed the *functional maps* framework, in which shape correspondence is modeled as a linear operator between spaces of functions on the shapes. The main advantage of functional maps is that finding correspondence boils down to a simple algebraic problem, as opposed to difficult combinatorial-type problems arising in, e.g., the computation of minimum-distortion maps. While several recent works showed that functional maps can be made resilient to missing parts or incomplete data [HWG14, KBBV15], overall this framework is not suitable for dealing with partial correspondence.

Contribution. In this paper, we propose an extension to the functional correspondence framework to allow dealing with partial correspondence. Specifically, we consider a scenario of matching a part of a deformed shape to some full model. Such scenarios are very common for instance in robotics applications, where one has to match an object acquired by means of a 3D scanner (and thus partially occluded) with a reference object known in advance. We use an explicit part model over which optimization is performed as in [BB08, BBBK09], as well as a regularization on the spectral representation of the functional correspondence accounting for a special structure of the Laplacian eigenfunctions as a result of part removal. Theoretical study of this behavior based on perturbation analysis of Laplacian matrices is another contribution of our work. We show experimentally that the proposed approach allows dealing with very challenging partial correspondence settings; further, we introduce a new benchmark to evaluate deformable partial correspondence methods, consisting of hundreds of shapes and ground-truth information.

The rest of the paper is organized as follows. In Section 2, we review the basic concepts in the spectral geometry and describe the functional correspondence approach. Section 3 studies the behavior of Laplacian eigenfunctions in the case of missing parts, motivating the regularizations used in the subsequent sections. Section 4 introduces our partial correspondence model, and Section 5 describes its implementation details. Section 6 presents experimental results, and finally, Section 7 concludes the paper.

2. Background

In this paper, we model shapes as compact connected 2-manifolds \mathcal{M} , possibly with boundary $\partial\mathcal{M}$. Given $f, g : \mathcal{M} \rightarrow \mathbb{R}$ some real scalar fields on the manifold, we define the standard inner product $\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} f(x)g(x)dx$, where integration is done using the area element induced by the Riemannian metric. We denote by $L^2(\mathcal{M}) = \{f : \mathcal{M} \rightarrow \mathbb{R} \mid \langle f, f \rangle_{\mathcal{M}} < \infty\}$ the space of square-integrable functions on \mathcal{M} .

The *intrinsic gradient* $\nabla_{\mathcal{M}}f$ and the positive semi-definite *Laplace-Beltrami operator* $\Delta_{\mathcal{M}}f = -\text{div}_{\mathcal{M}}(\nabla_{\mathcal{M}}f)$ generalize the notions of gradient and Laplacian to manifolds. The Laplace-Beltrami operator admits an eigen-decomposition

$$\Delta_{\mathcal{M}}\phi_i(x) = \lambda_i\phi_i(x) \quad x \in \text{int}(\mathcal{M}) \quad (1)$$

$$\langle \nabla_{\mathcal{M}}\phi_i(x), \hat{n}(x) \rangle = 0 \quad x \in \partial\mathcal{M}, \quad (2)$$

with homogeneous Neumann boundary conditions (2) if \mathcal{M} has a boundary (here \hat{n} denotes the normal vector to the boundary), where $0 = \lambda_1 < \lambda_2 \leq \dots$ are eigenvalues and ϕ_1, ϕ_2, \dots are the corresponding eigenfunctions (or eigenvectors). The eigenfunctions form an orthonormal basis on $L^2(\mathcal{M})$, i.e., $\langle \phi_i, \phi_j \rangle_{\mathcal{M}} = \delta_{ij}$, generalizing the classical

Fourier analysis: a function $f \in L^2(\mathcal{M})$ can be expanded into the *Fourier series* as

$$f(x) = \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \phi_i(x). \quad (3)$$

Functional correspondence. Let us be now given two manifolds, \mathcal{N} and \mathcal{M} . Ovsjanikov *et al.* [OBCS*12] proposed modeling *functional correspondence* between shapes as a linear operator $T : L^2(\mathcal{N}) \rightarrow L^2(\mathcal{M})$. One can easily see that classical vertex-wise correspondence is a particular setting where T maps delta-functions to delta-functions.

Assuming to be given two orthonormal bases $\{\phi_i\}_{i \geq 1}$ and $\{\psi_i\}_{i \geq 1}$ on $L^2(\mathcal{N})$ and $L^2(\mathcal{M})$ respectively, the functional correspondence can be expressed w.r.t. to these bases as follows:

$$\begin{aligned} Tf &= T \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{N}} \phi_i = \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{N}} T \phi_i \\ &= \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{N}} \underbrace{\langle T \phi_i, \psi_j \rangle_{\mathcal{M}}}_{c_{ij}} \psi_j, \end{aligned} \quad (4)$$

Thus, T amounts to a linear transformation of the Fourier coefficients of f from basis $\{\phi_i\}_{i \geq 1}$ to basis $\{\psi_i\}_{i \geq 1}$, which is captured by the coefficients c_{ij} . Truncating the Fourier series (4) at the first k coefficients, one obtains a rank- k approximation of T , represented in the bases $\{\phi_i, \psi_i\}_{i \geq 1}$ as a $k \times k$ matrix $\mathbf{C} = (c_{ij})$.

In order to compute \mathbf{C} , Ovsjanikov *et al.* [OBCS*12] assume to be given a set of q corresponding functions $\{f_1, \dots, f_q\} \subseteq L^2(\mathcal{N})$ and $\{g_1, \dots, g_q\} \subseteq L^2(\mathcal{M})$. Denoting by $a_{ij} = \langle f_j, \phi_i \rangle_{\mathcal{N}}$ and $b_{ij} = \langle g_j, \psi_i \rangle_{\mathcal{M}}$ the $k \times q$ matrices of the respective Fourier coefficients, functional correspondence boils down to the linear system

$$\mathbf{C}\mathbf{A} = \mathbf{B}. \quad (5)$$

If $q \geq k$, the system (5) is (over-)determined and is solved in the least squares sense to find \mathbf{C} .

Structure of \mathbf{C} . We note that the coefficients \mathbf{C} depend on the choice of the bases. In particular, it is convenient to use the eigenfunctions of the Laplace-Beltrami operators of \mathcal{N} and \mathcal{M} as the bases $\{\phi_i, \psi_i\}_{i \geq 1}$; truncating the series at the first k coefficients has the effect of ‘low-pass’ filtering thus producing smooth correspondences. In the following, this will be our tacit basis choice.

Furthermore, note that the system (5) has qk equations and k^2 variables. However, in many situations the actual number of variables is significantly smaller, as \mathbf{C} manifests a certain structure which can be taken advantage of. In particular, if \mathcal{N} and \mathcal{M} are isometric and have simple spectrum (*i.e.*, the Laplace-Beltrami eigenvalues have no multiplicity), then $T\phi_i = \pm\psi_i$, or in other words, $c_{ij} = \pm\delta_{ij}$. In more realistic scenarios (approximately isometric shapes), the matrix \mathbf{C} would manifest a funnel-shaped structure, with the majority of elements distant from the diagonal close to zero.

Discretization. In the discrete setting, the manifold \mathcal{N} is sampled at n points x_1, \dots, x_n which are connected by edges $E = E_i \cup E_b$ and faces F , forming a manifold triangular mesh (V, E, F) . We denote by E_i and E_b the interior and boundary edges respectively. A function on the manifold is represented by an n -dimensional vector $\mathbf{f} = (f(x_1), \dots, f(x_n))^T$. The discretization of the Laplacian takes the form of an $n \times n$ sparse matrix $\mathbf{L} = -\mathbf{S}^{-1}\mathbf{W}$ using the classical cotangent formula [Mac49, Duf59, PP93],

$$w_{ij} = \begin{cases} (\cot \alpha_{ij} + \cot \beta_{ij})/2 & ij \in E_i; \\ (\cot \alpha_{ij})/2 & ij \in E_b; \\ -\sum_{k \neq i} w_{ik} & i = j; \\ 0 & \text{else;} \end{cases} \quad (6)$$

where $\mathbf{S} = \text{diag}(s_1, \dots, s_n)$, $s_i = \frac{1}{3} \sum_{jk:ijk \in F} s_{ijk}$ denotes the local area element at vertex i , s_{ijk} denotes the area of triangle ijk , and α_{ij}, β_{ij} denote the angles $\angle ikj, \angle jhi$ of the triangles sharing the edge ij (see Fig. 2).

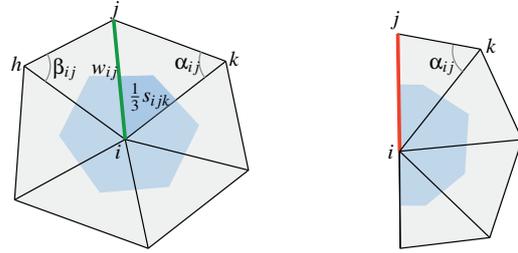


Figure 2: Discretization of the Laplace-Beltrami operator on a triangular mesh for interior edges (green, left) and boundary edges (red, right).

The first k eigenfunctions and eigenvalues of the Laplacian are computed by performing the generalized eigen-decomposition $\mathbf{W}\Phi = \mathbf{S}\Phi\Lambda$, where $\Phi = (\phi_1, \dots, \phi_k)$ is an $n \times k$ matrix containing as columns the discretized eigenfunctions and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_k)$ is the diagonal matrix of the corresponding eigenvalues. The computation of Fourier coefficients is performed by $\mathbf{a} = \Phi^T \mathbf{S}\mathbf{f}$.

3. Laplacian eigenvectors and eigenvalues under partiality

When one of the two shapes has missing parts, the assumption of approximate isometry does not hold anymore and a direct application of the method of Ovsjanikov *et al.* (*i.e.*, solving system (5)) would not produce meaningful results. However, as we show in this section, the matrix \mathbf{C} still exhibits a particular structure which can be exploited to drive the matching process.

We assume to be given a full shape \mathcal{M} and a part thereof $\mathcal{N} \subset \mathcal{M}$. We further denote by $\overline{\mathcal{N}} = \mathcal{M} \setminus \mathcal{N}$ the remaining vertices of \mathcal{M} . The manifolds \mathcal{M} and \mathcal{N} are discretized as triangular meshes with m and n vertices, respectively, and

$\bar{n} = m - n$. The scenario we consider in this paper concerns the problem of matching an approximately isometric deformation of part \mathcal{N} to the full shape \mathcal{M} (part-to-whole matching). Our goal is to characterize the eigenvalues and eigenvectors of the Laplacian $\mathbf{L}_{\mathcal{M}}$ in terms of perturbations of the eigenvalues and eigenvectors of the Laplacians $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\bar{\mathcal{N}}}$ [MH88]. We tacitly assume that homogeneous Neumann boundary conditions (2) apply.

3.1. Block-diagonal case

For the simplicity of analysis, let us first consider a simplified scenario in which \mathcal{N} and $\bar{\mathcal{N}}$ are *disconnected*, *i.e.*, there exist no links between the respective boundaries $\partial\mathcal{N}$ and $\partial\bar{\mathcal{N}}$. W.l.o.g., we can assume that the vertices in \mathcal{M} are ordered such that the vertices in \mathcal{N} come before those in $\bar{\mathcal{N}}$. With this ordering, the $m \times m$ Laplacian matrix $\mathbf{L}_{\mathcal{M}}$ is block-diagonal, with an $n \times n$ block $\mathbf{L}_{\mathcal{N}}$ and an $\bar{n} \times \bar{n}$ block $\mathbf{L}_{\bar{\mathcal{N}}}$. The (sorted) eigenvalues of $\mathbf{L}_{\mathcal{M}}$ form a mixed sequence composed of the eigenvalues of $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\bar{\mathcal{N}}}$. Similarly, the eigenvectors of $\mathbf{L}_{\mathcal{M}}$ correspond to the eigenvectors of the two sub-matrices, zero-padded to the correct size (Fig. 3).

Structure of \mathbf{C} under partiality. Suppose we are given the first k Laplace-Beltrami eigenvalues of the full shape \mathcal{M} and of its part \mathcal{N} . Since the spectrum of $\mathbf{L}_{\mathcal{M}}$ is an interleaved sequence of the eigenvalues of $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\bar{\mathcal{N}}}$, only the first $r < k$ eigenvalues of $\mathbf{L}_{\mathcal{N}}$ will appear among the first k eigenvalues of $\mathbf{L}_{\mathcal{M}}$. The remaining $k - r$ eigenvalues of $\mathbf{L}_{\mathcal{N}}$ will only appear further along the spectrum of $\mathbf{L}_{\mathcal{M}}$ (see Fig. 6 for an example where $k = 50$ and $r = 21$). The same argument holds for the associated eigenfunctions, as illustrated in Fig. 4: if ϕ_i is an eigenvector of $\mathbf{L}_{\mathcal{N}}$, then $\mathbf{L}_{\mathcal{M}}$ also has an eigenvector ψ_j such that $\phi_i = \mathbf{T}\psi_j$, where $\mathbf{T} = (\mathbf{I}_{n \times n}, \mathbf{0})^\top$ and $i < j$.

This analysis leads us to the following simple observation: the partial functional map between \mathcal{N} and \mathcal{M} is represented

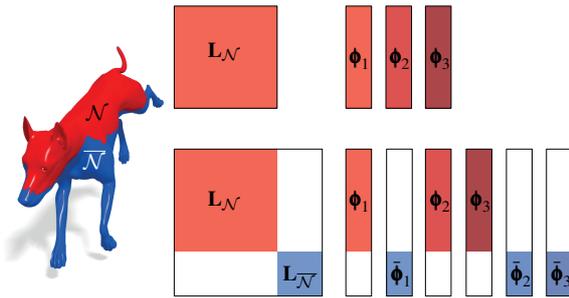


Figure 3: The eigenvalues and eigenvectors of a block-diagonal Laplacian $\mathbf{L}_{\mathcal{M}}$ are an interleaved sequence of the eigenpairs from the two blocks $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\bar{\mathcal{N}}}$.

in the spectral domain by the matrix of inner products $c_{ij} = \langle \mathbf{T}\phi_i, \psi_j \rangle_{\mathcal{M}}$, which has a *slanted-diagonal* structure with a slope r/k (see examples in Figs. 1, 4 where this structure is manifested approximately). Consequently, the last $k - r$ columns of matrix \mathbf{C} are zero such that $r = \text{rank}(\mathbf{C})$. The value r can be estimated by simply comparing the spectra of the two shapes, as shown in Fig. 6. Note that this behavior is consistent with Weyl's asymptotic law [Wey11], according to which the Laplacian eigenvalues grow linearly, with rate inversely proportional to surface area.

3.2. Perturbation analysis

We will now show that these properties still approximately hold when the Laplacian matrix $\mathbf{L}_{\mathcal{M}}$ is not perfectly block-diagonal, *i.e.*, when \mathcal{N} and $\bar{\mathcal{N}}$ are joined along their boundaries. Roughly speaking, the main observation is that in this case as well the matrix \mathbf{C} has a slanted diagonal structure, where the diagonal angle depends on the relative area of the part, and the diagonal 'sharpness' depends on the position and length of the cut.

Here we assume w.l.o.g. that within \mathcal{N} the boundary vertices $\partial\mathcal{N}$ are indexed at the end, while within $\bar{\mathcal{N}}$ the boundary vertices $\partial\bar{\mathcal{N}}$ are indexed at the beginning. Then, there is a boundary band $\mathcal{B} = \partial\mathcal{N} \cup \partial\bar{\mathcal{N}}$ such that only the entries of the Laplacians $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\bar{\mathcal{N}}}$ between vertices in \mathcal{B} are affected by the cut (Fig. 5).

We define the parametric matrix

$$\mathbf{L}(t) = \left(\begin{array}{c|c} \mathbf{L}_{\mathcal{N}} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{L}_{\bar{\mathcal{N}}} \end{array} \right) + t \left(\begin{array}{c|c} \mathbf{P}_{\mathcal{N}} & \mathbf{P} \\ \hline \mathbf{P}^\top & \mathbf{P}_{\bar{\mathcal{N}}} \end{array} \right), \quad (7)$$

where

$$\mathbf{P}_{\mathcal{N}} = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_{\mathcal{N}} \end{pmatrix}, \quad \mathbf{P}_{\bar{\mathcal{N}}} = \begin{pmatrix} \mathbf{D}_{\bar{\mathcal{N}}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad \mathbf{P} = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{E} & \mathbf{0} \end{pmatrix}$$

are matrices of size $n \times n$, $\bar{n} \times \bar{n}$, and $n \times \bar{n}$, respectively. Here $\mathbf{D}_{\mathcal{N}}$ and $\mathbf{D}_{\bar{\mathcal{N}}}$ represent the variations of the Laplacians $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\bar{\mathcal{N}}}$ within nodes in $\partial\mathcal{N}$ and $\partial\bar{\mathcal{N}}$ respectively, while \mathbf{E} represents the variations across the boundary. The parameter t is such that $\mathbf{L}(1) = \mathbf{L}_{\mathcal{M}}$, while for $t = 0$ we get back

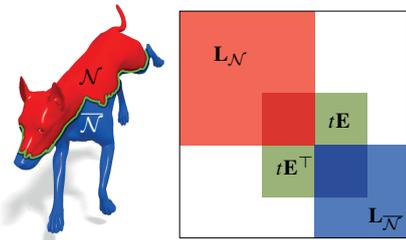


Figure 5: The matrix $\mathbf{L}(t)$ is obtained as a perturbation of the block-diagonal Laplacian in the boundary band (shown in green).

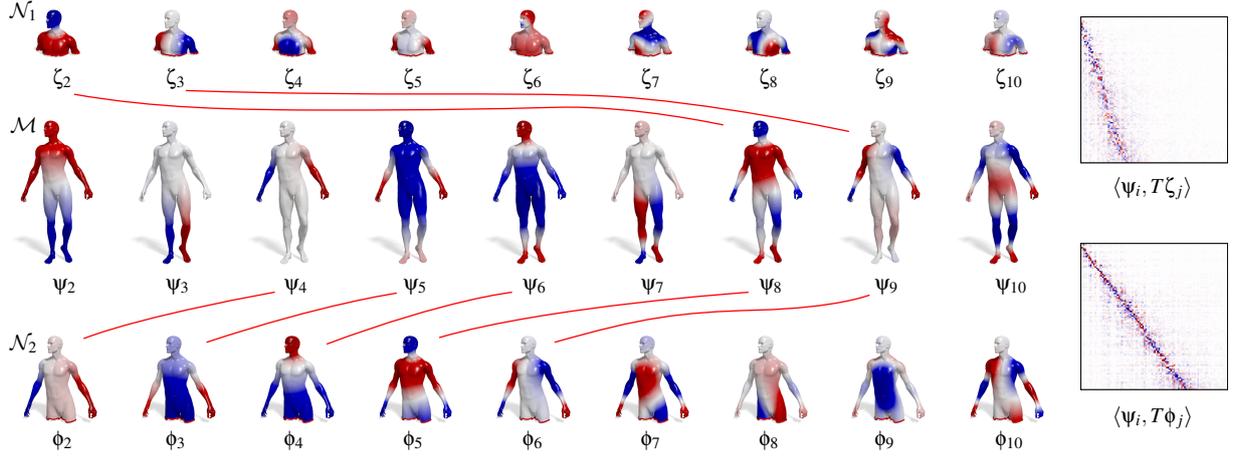


Figure 4: First ten eigenfunctions of a full shape \mathcal{M} and two parts $\mathcal{N}_1, \mathcal{N}_2 \subset \mathcal{M}$ with different surface area. All eigenfunctions of the partial shapes have a corresponding eigenfunction Ψ_i on the full shape for some i ; the correspondence between eigenfunctions follows from the correspondence between eigenvalues (see also Fig. 6). This is reflected in functional maps with different diagonal slopes, where the slope depends on the area ratios of the two surfaces (by Weyl’s law).

to the disconnected case of Fig. 3. In what follows, we perform a differential analysis at $t = 0$ and therefore analyze the change in eigenvalues and eigenvectors of $\mathbf{L}(t)$ as we interpolate between the seen (\mathcal{N}) and unseen ($\bar{\mathcal{N}}$) parts for $t = 0$, to the full shape \mathcal{M} for $t = 1$.

Note that with the appropriate ordering of the vertices, the matrices $\mathbf{D}_{\mathcal{N}}$ and $\mathbf{D}_{\bar{\mathcal{N}}}$ will have a band-diagonal structure. In fact, a cut through an edge will affect the values of the discrete Laplacian matrix \mathbf{L} only at the entries corresponding to the vertices at the extremities of the edge, and to the edges laying in the same triangle as the cut edge. For example, looking at Fig. 2, a cut through edge (i, j) will affect the diagonal entries l_{ii} and l_{jj} as well as the off-diagonal entries l_{ih} , l_{ik} , l_{jh} , and l_{jk} . Note also that the continuity of the cut implies that two of the four off-diagonal entries will be cut as well, leaving no more than two affected edges on any side of the cut. As a result, the entries of the Laplacian affected by the cut correspond to the nodes and edges in a path along the boundary of the cut.

We further note that although here we consider the cotangent Laplacian for simplicity of analysis, similar results hold for Laplacians that are not strictly local, but locally dominant [BSW09b], *i.e.*, most of their L_2 norm is due to the elements in a tight boundary layer.

Theorem 1 Let $\mathbf{L}_{\mathcal{N}} + t\mathbf{P}_{\mathcal{N}} = \Phi(t)^\top \mathbf{\Lambda}(t)\Phi(t)$, where $\mathbf{\Lambda}(t) = \text{diag}(\lambda_1(t), \dots, \lambda_n(t))$ is a diagonal matrix of eigenvalues, and $\Phi(t)$ are the corresponding eigenvectors. The derivative of the non-trivial eigenvalues is given by

$$\frac{d}{dt}\lambda_i = \sum_{v,w \in \partial\mathcal{N}} (\mathbf{P}_{\mathcal{N}})_{vw} \phi_{iv} \phi_{iw} = \Phi_i^\top \mathbf{P}_{\mathcal{N}} \Phi_i. \quad (8)$$

Proof: See Appendix C.

Theorem 1 establishes that the (first-order) change in the eigenvalues of the partial shape \mathcal{N} only depends on the change in the Dirichlet energy of the corresponding eigenvectors along the boundary $\partial\mathcal{N}$ (recall that for eigenvector Φ_i , the Dirichlet energy is defined as $\Phi_i^\top (\mathbf{L}_{\mathcal{N}} + t\mathbf{P}_{\mathcal{N}})\Phi_i = \Phi_i^\top \mathbf{L}_{\mathcal{N}}\Phi_i + t\Phi_i^\top \mathbf{P}_{\mathcal{N}}\Phi_i$). This means that the eigenvalues are perturbed depending on the *length* and *position* of the cut. Note that the estimate given in Eq. (8) can not typically be computed directly, as this would assume knowledge of the correspondence between \mathcal{N} and \mathcal{M} . However, by virtue of this result, we can establish approximate correspondence between the eigenvalues $\lambda_i^{\mathcal{N}}$ of $\mathbf{L}_{\mathcal{N}}$ and a subset of the eigenvalues $\lambda_j^{\mathcal{M}}$ of $\mathbf{L}_{\mathcal{M}}$ (which are now not exactly equal as in the block-diagonal case). We do this in order to estimate the slope of \mathbf{C} . Specifically, we compute

$$r = \max\{i \mid \lambda_i^{\mathcal{N}} < \max_{j=1}^k \lambda_j^{\mathcal{M}}\}; \quad (9)$$

the slope of \mathbf{C} can now be estimated as r/k , as explained in Section 3.1 (see also Fig. 6 for a visual illustration of the estimation of r).

Theorem 2 Assume that $\mathbf{L}_{\mathcal{N}}$ has distinct eigenvalues ($\lambda_i \neq \lambda_j$ for $i \neq j$), and furthermore, the non-zero eigenvalues are all distinct from the eigenvalues of $\mathbf{L}_{\bar{\mathcal{N}}}$ ($\lambda_i \neq \bar{\lambda}_j$ for all i, j). Let $\mathbf{L}_{\mathcal{N}} + t\mathbf{P}_{\mathcal{N}} = \Phi(t)^\top \mathbf{\Lambda}(t)\Phi(t)$, where $\mathbf{\Lambda}(t) = \text{diag}(\lambda_1(t), \dots, \lambda_n(t))$ is a diagonal matrix of eigenvalues, and $\Phi(t)$ are the corresponding eigenvectors. Then, the derivative of the non-constant eigenvector is given by

$$\frac{d}{dt}\Phi_i = \sum_{\substack{j=1 \\ j \neq i}}^n \frac{\Phi_i^\top \mathbf{P}_{\mathcal{N}} \Phi_j}{\lambda_i - \lambda_j} \Phi_j + \sum_{j=1}^{\bar{n}} \frac{\Phi_i^\top \mathbf{P} \bar{\Phi}_j}{\lambda_i - \bar{\lambda}_j} \bar{\Phi}_j. \quad (10)$$

Proof: See Appendix C.

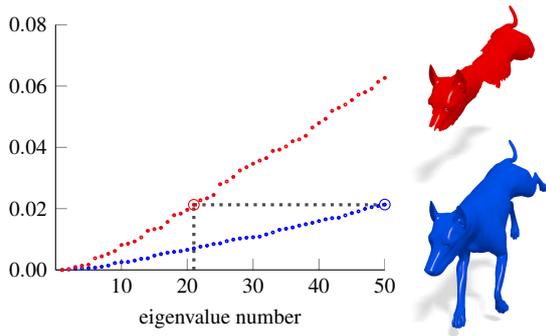


Figure 6: Neumann spectra of a full shape and a part of it. The eigenvalues of the partial shape (in red) are approximately preserved under the partiality transformation (see Theorem 1), and appear perturbed in the spectrum of the full shape (in blue). This simple observation allows us to estimate the diagonal slope of the functional map relating the two shapes; in this example, the slope is equal to 21/50.

Remark. If $\mathbf{L}_{\overline{\mathcal{N}}}$ shares some eigenvalues with $\mathbf{L}_{\mathcal{N}}$, the second sum in (10) would be slightly different [MH88], but would still only have support over $\overline{\mathcal{N}}$.

We conclude from Theorem 2 that the perturbation associated with the partiality transformation gives rise to a mixing of eigenspaces. The second summation in (10) has support over $\overline{\mathcal{N}}$ and thus provides the completion of the eigenfunction on the missing part. The first summation in (10) is responsible for the modifications of the eigenvectors over the nodes in \mathcal{N} . Here the numerator has a term $\phi_i^\top \mathbf{P}_{\mathcal{N}} \phi_j$ which, since $\mathbf{D}_{\mathcal{N}}$ is band-diagonal and diagonally dominant, acts as a dot product of the eigenvectors over the boundary band. This points to large mixing of eigenvectors with a strong co-presence near the boundary. In turn, the term $\lambda_i - \lambda_j$ at the denominator forces a strong mixing of eigenvectors corresponding to similar eigenvalues. This results in an amplification of the variation for higher eigenvalues, as eigenvalues tend to densify on the higher end of the spectrum, and explains the funnel-shaped spread of the matrix \mathbf{C} visible at high frequencies (see Fig. 8).

Similarly to the case of eigenvalues, the eigenvectors are also perturbed depending on the length and position of the cut. The variation of the eigenvectors due to the mixing within the partial shape can be reduced either by shortening the boundary of the cut, or by reducing the strength of the boundary interaction. The latter can be achieved by selecting a boundary along which eigenvectors with similar eigenvalues are either orthogonal, or both small. The *boundary interaction strength* can be quantified by considering the following function (we refer to Appendix C for a derivation):

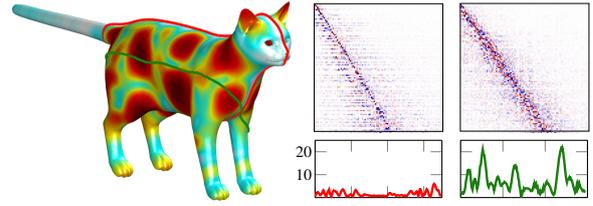


Figure 7: Left: A model is cut in two different ways (red and green curves) with cuts of same length. The off-diagonal dispersion depends mainly on the position of each cut. Function f (11) is plotted over the model. Middle: Ground-truth functional map between the complete model and the partial shape produced by the red cut (top), and values of f along the cut (bottom). Right: Plots associated to the green cut.

$$f(\mathbf{v}) = \sum_{\substack{i,j=1 \\ j \neq i}}^n \left(\frac{\phi_{iv} \phi_{jv}}{\lambda_i - \lambda_j} \right)^2. \quad (11)$$

Fig. 7 shows an example of two different cuts with different interaction strengths, where the function f is plotted on top of the cat model. The cuts plotted in the figure have the same length, but one cut goes along a symmetry axis of the shape and through low values of f , while the other goes through rather high values of f . This is manifested in the dispersion of the slanted diagonal structure of the matrix \mathbf{C} (larger in the second case).

4. Partial functional maps

As stated before, throughout the paper we consider the setting where we are given a full model shape \mathcal{M} and another query shape \mathcal{N} that corresponds to an approximately isometrically deformed part $\mathcal{M}' \subset \mathcal{M}$.

Following [BB08], we model the part \mathcal{M}' by means of an indicator function $v: \mathcal{M} \rightarrow \{0, 1\}$ such that $v(x) = 1$ if $x \in \mathcal{M}'$ and zero otherwise. Assuming that v is known, the *partial functional correspondence* between \mathcal{N} and \mathcal{M} can be expressed as $Tf = vg$, where v can be regarded as a kind of mask, and anything outside the region where $v = 1$ should be ignored. Expressed w.r.t. bases $\{\phi_i\}_{i \geq 1}$ and $\{\psi_i\}_{i \geq 1}$, the partial functional correspondence takes the form $\mathbf{C}\mathbf{A} = \mathbf{B}(v)$, where $\mathbf{B}(v)$ denotes a matrix of weighted inner products with elements given by $b_{ij}(v) = \int_{\mathcal{M}} v(x) \psi_i(x) g_j(x) dx$ (when $v(x) \equiv 1$, \mathbf{B} is simply the matrix of Fourier coefficients defined in (4)).

This brings us to the problem we are considering throughout this paper, involving optimization w.r.t. correspondence (encoded by the coefficients \mathbf{C}) and the part v ,

$$\min_{\mathbf{C}, v} \|\mathbf{C}\mathbf{A} - \mathbf{B}(\eta(v))\|_{2,1} + \rho_{\text{corr}}(\mathbf{C}) + \rho_{\text{part}}(v), \quad (12)$$

where $\eta(t) = \frac{1}{2} (\tanh(2t - 1) + 1)$ saturates the part indica-

tor function between zero and one (see below). Here ρ_{corr} and ρ_{part} denote regularization terms for the correspondence and the part, respectively; these terms are explained below. We use the $L_{2,1}$ matrix norm (equal to the sum of L_2 -norms of matrix columns) to handle possible outliers in the corresponding data, as such a norm promotes column-sparse matrices. A similar norm was adopted in [HWG14] to handle spurious maps in shape collections.

Note that in order to avoid a combinatorial optimization over binary-valued v , we use a continuous v with values in the range $(-\infty, +\infty)$, saturated by the non-linearity η . This way, $\eta(v)$ becomes a soft membership function with values in the range $[0, 1]$.

Part regularization. Similarly to [BB08, PBB13], we try to find the part with area closest to that of the query and with shortest boundary. This can be expressed as

$$\begin{aligned} \rho_{\text{part}}(v) = & \mu_1 \left(\text{area}(\mathcal{N}) - \int_{\mathcal{M}} \eta(v) dx \right)^2 \quad (13) \\ & + \mu_2 \int_{\mathcal{M}} \xi(v) \|\nabla_{\mathcal{M}} v\| dx, \end{aligned}$$

where $\xi(t) \approx \delta\left(\eta(t) - \frac{1}{2}\right)$ and the norm is on the tangent space. The μ_2 -term in (13) is an intrinsic version of the Mumford-Shah functional [MS89], measuring the length of the boundary of a part represented by a (soft) membership function. This functional was used previously in image segmentation applications [VC02].

Correspondence regularization. For the correspondence, we use the penalty

$$\begin{aligned} \rho_{\text{corr}}(\mathbf{C}) = & \mu_3 \|\mathbf{C} \circ \mathbf{W}\|_F^2 + \mu_4 \sum_{i \neq j} (\mathbf{C}^\top \mathbf{C})_{ij}^2 \\ & + \mu_5 \sum_i ((\mathbf{C}^\top \mathbf{C})_{ii} - d_i)^2, \quad (14) \end{aligned}$$

where \circ denotes Hadamard (element-wise) matrix product. The μ_3 -term models the special slanted-diagonal structure of \mathbf{C} that we observe in partial matching problems (see Fig. 8); the theoretical motivation for this behavior was presented in Sec. 3. Here, \mathbf{W} is a weight matrix with zeros along the slanted diagonal and large values outside (see Fig. 8; details on the computation of \mathbf{W} are provided in Appendix A).

The μ_4 -term promotes orthogonality of \mathbf{C} by penalizing the off-diagonal elements of $\mathbf{C}^\top \mathbf{C}$. The reason is that for isometric shapes, the functional map is volume-preserving, and this is manifested in orthogonal \mathbf{C} [OBCS*12]. Note that differently from the classical case (*i.e.*, full shapes), in our setting we can only require area preservation going in the direction from partial to complete model, as also expressed by the μ_1 -term in (13). For this reason, we do not impose any restrictions on $\mathbf{C}\mathbf{C}^\top$ and we say that the matrix is *semi*-orthogonal.

Finally, note that due to the low-rank nature of \mathbf{C} we can

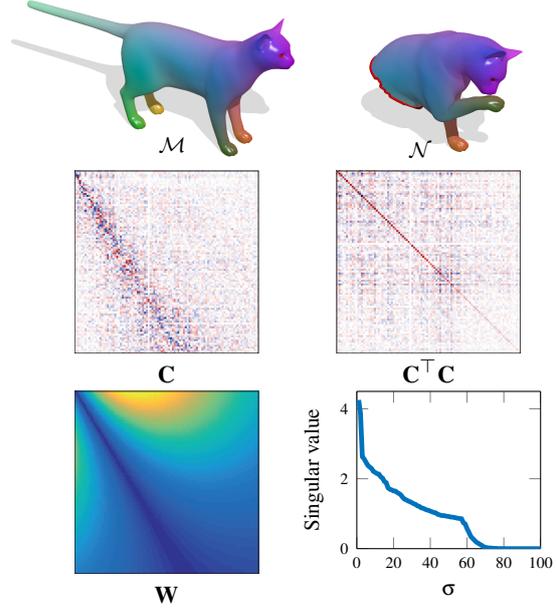


Figure 8: A partial functional map \mathbf{C} from \mathcal{N} to \mathcal{M} has a slanted-diagonal structure (second row, left). The low-rank nature of such a map is also manifested in its singular values (bottom right). If the map is volume-preserving, then its full-rank sub-matrix is orthogonal: observe how $\mathbf{C}^\top \mathbf{C}$ approximates the identity, with a trail of small values along the diagonal corresponding to the almost-zero block of \mathbf{C} .

not expect the product $\mathbf{C}^\top \mathbf{C}$ to be full rank. Indeed, we expect elements off the slanted diagonal of \mathbf{C} to be close to zero and thus $\mathbf{C}^\top \mathbf{C} \approx \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$. The μ_5 -term in (14) models this behavior, where vector $\mathbf{d} = (d_1, \dots, d_k)$ determines how many singular values of \mathbf{C} are non-zero (the estimation of \mathbf{d} is straightforward, and described in Appendix A).

Remark. The fact that matrix \mathbf{C} is low-rank is a direct consequence of partiality. This can be understood by recalling from Eq. (4) that the (non-truncated) functional map representation amounts to an orthogonal change of basis; since in the standard basis the correspondence matrix is low-rank (as it contains zero-sum rows), this property is preserved by the change of basis.

In Fig. 8 we show an example of a ground-truth partial functional map \mathbf{C} , illustrating its main properties.

Alternating scheme. To solve the optimization problem (12), we perform an alternating optimization w.r.t. to \mathbf{C} and v , repeating the following steps until convergence:

C-step: Fix v^* , solve for correspondence \mathbf{C}

$$\min_{\mathbf{C}} \|\mathbf{C}\mathbf{A} - \mathbf{B}(\eta(v^*))\|_{2,1} + \rho_{\text{corr}}(\mathbf{C}). \quad (15)$$

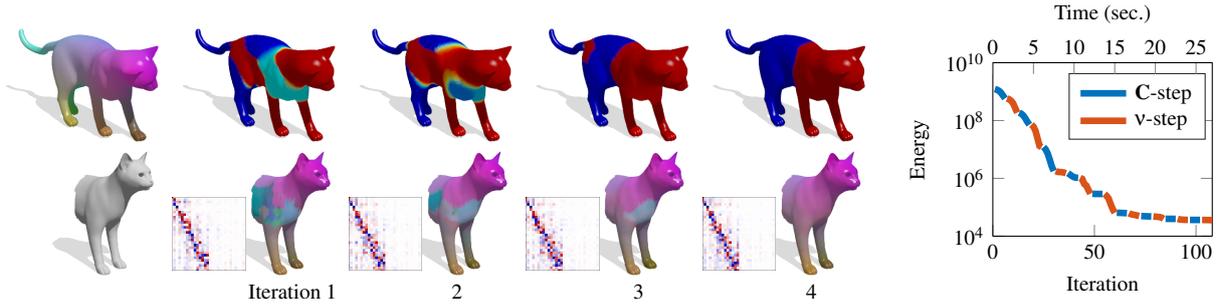


Figure 9: An example of the matching process operating on two shapes from TOSCA. The algorithm alternately optimizes over corresponding part (top row) and functional correspondence (bottom row). Corresponding points between full and partial shape are shown with the same color. This solution was obtained by using 30 eigenfunctions on both manifolds.

V-step: Fix \mathbf{C}^* , solve for part v

$$\min_v \|\mathbf{C}^* \mathbf{A} - \mathbf{B}(\eta(v))\|_{2,1} + \rho_{\text{part}}(v). \quad (16)$$

A practical example of the alternating scheme applied to a pair of shapes is shown in Fig. 9.

5. Implementation

We refer to Appendix A for the discretization of the regularization terms appearing in (13) and (14).

Numerical optimization. We implemented our matching framework in Matlab/C++ using the manifold optimization toolbox [BMAS14]. Each optimization step was performed by the method of nonlinear conjugate gradients. Detailed derivations of the involved gradients can be found in Appendix B. We initialize the alternating scheme by fixing $\mathbf{v}^* = \mathbf{1}$ (a vector of m ones), $\mathbf{C} = \mathbf{W}$, and by optimizing over \mathbf{C} . In all our experiments we observed convergence in 3-5 outer iterations (around 5 mins. for a pair of shapes).

Refinement. In order to account for noisy data, we also run a refinement step after each C-step. Specifically, assume \mathbf{C}^* is a local optimum of problem (15), and consider the term $\|\mathbf{C}^* \Phi^T - \Psi^T \Pi\|_F$, where Π is a left-stochastic binary matrix assigning each column of Ψ^T to the nearest column of $\mathbf{C}^* \Phi^T$; this is done by n nearest-neighbor searches in \mathbb{R}^k , one per column. Given the optimal Π^* , we solve for the map \mathbf{C} minimizing $\|\mathbf{C} \Phi^T - \Psi^T \Pi^*\|_F$ plus the μ_4, μ_5 terms of Eq. (14). We alternate the \mathbf{C}^* and Π^* steps until convergence. This refinement step can be seen as a generalization to partial maps of the ICP-like technique found in [OBBS*12], and can be interpreted as an attempt to improve the alignment between the spectral embeddings of the two shapes. Further note that matrix Π^* encodes the point-wise correspondence between \mathcal{N} and \mathcal{M} , which is used to evaluate the accuracy of our method.

6. Experimental results

Datasets. As base models, we use shapes from the TOSCA dataset [BBK08], consisting of 76 nearly-isometric shapes subdivided into 8 classes. Each class comes with a “null” shape in a standard pose (extrinsically bilaterally symmetric), and ground-truth correspondences are provided for all shapes within the same class. In order to make the datasets more challenging and avoid compatible triangulations, all shapes were remeshed to 10K vertices by iterative pair contractions [GH97]. Then, missing parts were introduced in the following ways:[†]

Regular cuts. The null shape of each class was cut with a plane at 6 different orientations, including an exact cut along the symmetry plane. The six cuts were then transferred to the remaining poses using the ground-truth correspondence, resulting in 456 partial shapes in total. Some examples are shown in Fig. 6 and 8.

Irregular holes. Given a shape and an “area budget” determining the fraction of area to keep (40%, 70%, and 90%), we produced additional shapes by an erosion process applied to the surface. Specifically, seed holes were placed at 5, 25, and 50 farthest samples over the shape; the holes were then enlarged to meet the specified area budget. The total number of shapes produced this way was 684. Examples of this dataset are shown in Fig. 11 and 13.

Range images. We simulated range images by taking orthographic projections of the original TOSCA shapes from different viewpoints. Each range image was produced via ray casting from a regular grid with a resolution of 100×150 pixels. Examples are shown in Fig. 13.

Point clouds. Point clouds were generated by taking a subset of shapes from the first two datasets. Each partial shape

[†] The datasets together with code for our method are available for download at <http://vision.in.tum.de/data/datasets/partial>.

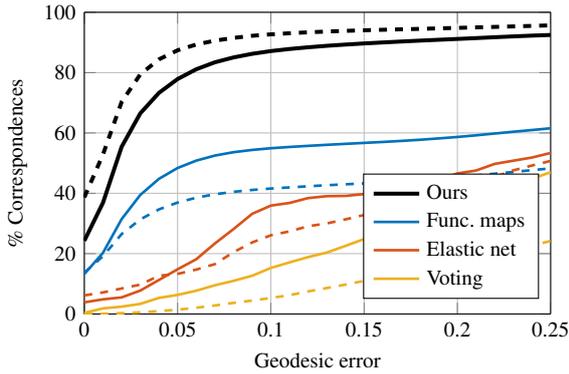


Figure 10: Correspondence quality of different methods evaluated using the Princeton protocol on partial TOSCA shapes with regular cuts (solid) and irregular holes (dotted).

was then resampled uniformly to 1000 farthest points, and the tessellation removed. See Fig. 13 for examples.

Where not specified otherwise, we use 120 random partial shapes for the first dataset and 80 for the second, equally distributed among the different classes. Each partial shape is then matched to the null shape of the corresponding class.

Error measure. For the evaluation of the correspondence quality, we used the Princeton benchmark protocol [KLF11] for point-wise maps. Assume that a correspondence algorithm produces a pair of points $(x, y) \in \mathcal{N} \times \mathcal{M}$, whereas the ground-truth correspondence is (x, y^*) . Then, the inaccuracy of the correspondence is measured as

$$\varepsilon(x) = \frac{d_{\mathcal{M}}(y, y^*)}{\text{area}(\mathcal{M})^{1/2}}, \quad (17)$$

and has units of normalized length on \mathcal{M} (ideally, zero). Here $d_{\mathcal{M}}$ is the geodesic distance on \mathcal{M} . The value $\varepsilon(x)$ is averaged over all shapes \mathcal{N} . We plot cumulative curves showing the percent of matches which have error smaller than a variable threshold.

Methods. We compared the proposed method with (full) functional maps [OBCS*12], elastic net [RTH*13], and the voting method of [SY14] using the code provided by the respective authors.

Local descriptors. Due to the particular nature of the problem, in all our experiments we only make use of dense, local descriptors as a data term. This is in contrast with the more common scenario in which full shapes are being matched – thus allowing to employ more robust, globally-aware features such as landmark matches, repeatable surface regions, and various spectral quantities [OBCS*12]. In our experiments, we used the extrinsic SHOT [TSDS10] descriptor, computed using 10 normal bins (352 dimensions in total).

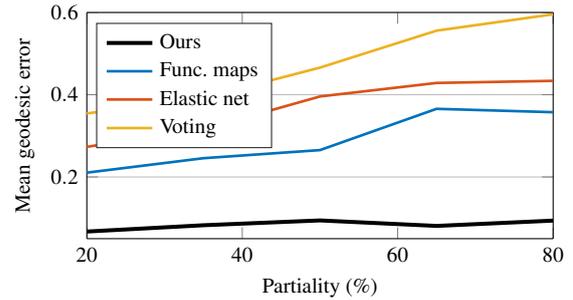


Figure 11: Correspondence quality (in terms of mean geodesic error, in % of diameter) obtained by different methods at increasing levels of partiality. Other methods show significant performance drop with increasing partiality, while the performance of our method is nearly-constant.

As opposed to [ART15, PBB13] which ignore points close to the boundary in order to avoid boundary effects, in our formulation we retain all shape points.

6.1. Sensitivity analysis

We conducted a set of experiments aimed at evaluating the sensitivity of our approach to different parametrizations. In order to reduce overfitting we only used a subset of TOSCA (regular cuts), namely composed of the *cat* and *victoria* shape classes (20 pairs).

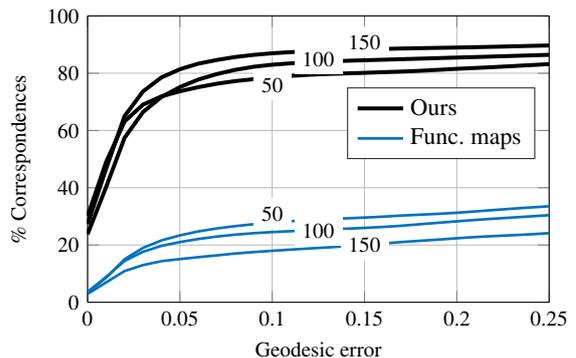


Figure 12: Correspondence quality obtained on a subset of TOSCA at increasing rank (reported as labels on top of the curves). Note the opposite behavior of the baseline approach and our regularized partial matching.



Figure 13: Examples of partial functional correspondence obtained with our method on meshes and point clouds from the proposed datasets. Notice how regions close to the boundary are still accurately matched despite the noisy descriptors.

Rank. In the first experiment we study the change in accuracy as the rank of the functional map is increased; this corresponds to using an increasing number of basis functions for the two shapes being matched. For this experiment we compare with the baseline method of Ovsjanikov *et al.* [OBCS*12] by using the same dense descriptors as ours. For fair comparisons, we did not impose map orthogonality or operator commutativity constraints [OBCS*12], which cannot obviously be satisfied due to partiality. The results of this experiment are reported in Fig. 12. As we can see from the plots, our method allows to obtain more accurate solutions as the rank increases, while an opposite behavior is observed for the other method.

Representation. Our method is general enough to be applied to different shape representations, as long as a proper discretization of the Laplace operator is available. In Fig. 13 we show some qualitative examples of correspondences produced by our algorithm on simulated point clouds and depth maps. Here we use the method described in [BSW09a] to

construct a discrete Laplacian on the point clouds. This is traditionally considered a particularly challenging problem in robotics and vision applications, with few methods currently capable of giving satisfactory solutions without exploiting controlled conditions or domain-specific information (*e.g.*, the knowledge that the shape being matched is that of a human). These are, to the best of our knowledge, the best results to be published so far for this class of problems.

6.2. Comparisons

We compared our method on the *cuts* and *holes* datasets (200 shape pairs in total); the results are shown in Fig. 10. As an additional experiment, we ran comparisons against [OBCS*12] across increasing amounts of partiality. The rationale behind this experiment is to show that, at little or no partiality, our approach converges to the one described in [OBCS*12], currently among the state of the art in non-rigid shape matching. However, as partiality increases so

does the sensitivity of the latter method. Fig. 11 shows the results of this experiment.

Parameters for our method were chosen on the basis of the sensitivity analysis. Specifically, we used $k = 100$ eigenfunctions per shape, and set $\mu_1 = \mu_3 = 1$, $\mu_4 = \mu_5 = 10^3$, and $\mu_2 = 10^2$. The different orders of magnitude for the μ coefficients are due to the fact that the regularizing terms operate at different scales. We also experimented with other values, but in all our tests we did not observe significant changes in accuracy. Additional examples of partial matchings obtained with our method are shown in Fig. 13.

7. Discussion and conclusions

In this paper we tackled the problem of dense matching of deformable shapes under partiality transformations. We cast our formulation within the framework of functional maps, which we adapted and extended to deal with this more challenging scenario. Our approach is fully automatic and makes exclusive use of dense local features as a measure of similarity. Coupled with a robust prior on the functional correspondence derived from a perturbation analysis of the shape Laplacians, this allowed us to devise an effective optimization process with remarkable results on very challenging cases. In addition to our framework for partial functional correspondence we also introduced two new datasets comprising hundreds of shapes, which we hope will foster further research on this challenging problem.

One of the main issues of our method concerns the existence of multiple optima, which is in turn related to the presence of non-trivial self-isometries on the considered manifolds. Since most natural shapes are endowed with intrinsic symmetries, one may leverage this knowledge in order to avoid inconsistent matchings. For example, including a smoothness prior on the correspondence might alleviate such imperfections and thus provide better-behaved solutions. Secondly, since the main focus of this paper is on tackling partiality rather than general deformations, our current formulation does not explicitly address the cases of topological changes and inter-class similarity (e.g., matching a man to a gorilla). However, the method can be easily extended to employ more robust descriptors such as pairwise features [RABT13, vKZH13], or to simultaneously optimize over ad-hoc functional bases on top of the correspondence. Finally, extending our approach to tackle entire shape collections, as opposed to individual pairs of shapes, represents a further exciting direction of research.

Acknowledgments

The authors thank Matthias Vestner, Vladlen Koltun, Aneta Stevanović, Zorah Lähler, Maks Ovsjanikov, and Ron Kimmel for useful discussions. ER is supported by an Alexander von Humboldt Fellowship. MB is partially supported by the ERC Starting Grant No. 307047 (COMET).

References

- [AMCO08] AIGER D., MITRA N. J., COHEN-OR D.: 4-points congruent sets for robust pairwise surface registration. *TOG* 27, 3 (2008), 85. 1
- [ART15] ALBARELLI A., RODOLÀ E., TORSELLO A.: Fast and accurate surface alignment through an isometry-enforcing game. *Pattern Recognition* 48 (2015), 2209–2226. 1, 9
- [BB08] BRONSTEIN A. M., BRONSTEIN M. M.: Not only size matters: regularized partial matching of nonrigid shapes. In *Proc. NORDIA* (2008). 1, 2, 6, 7
- [BBK09] BRONSTEIN A., BRONSTEIN M., BRUCKSTEIN A., KIMMEL R.: Partial similarity of objects, or how to compare a centaur to a horse. *IJCV* 84, 2 (2009), 163–183. 1, 2
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103, 5 (2006), 1168–1172. 1
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 8
- [BMAS14] BOUMAL N., MISHRA B., ABSIL P.-A., SEPULCHRE R.: Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research* 15 (2014), 1455–1459. URL: <http://www.manopt.org>. 8
- [BSW09a] BELKIN M., SUN J., WANG Y.: Constructing laplace operator from point clouds in rd. In *Proc. SODA* (2009), Society for Industrial and Applied Mathematics, pp. 1031–1040. 10
- [BSW09b] BELKIN M., SUN J., WANG Y.: Discrete laplace operator for meshed surfaces. In *Proc. SODA*. 2009, pp. 1031–1040. 5
- [BWW*14] BRUNTON A., WAND M., WUHRER S., SEIDEL H.-P., WEINKAUF T.: A low-dimensional representation for robust partial isometric correspondences computation. *Graphical Models* 76, 2 (2014), 70 – 85. 2
- [CGH14] CHEN Y., GUIBAS L., HUANG Q.: Near-optimal joint object matching via convex relaxation. In *Proc. ICML* (2014), pp. 100–108. 2
- [CRA*16] COSMO L., RODOLÀ E., ALBARELLI A., MÉMOLI F., CREMERS D.: Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum* (2016). 2
- [Duf59] DUFFIN R. J.: Distributed and lumped networks. *Journal of Mathematics and Mechanics* 8, 5 (1959), 793–826. 3
- [GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proc. SIGGRAPH* (1997), pp. 209–216. 8
- [HWG14] HUANG Q., WANG F., GUIBAS L. J.: Functional map networks for analyzing and exploring large shape collections. *TOG* 33, 4 (2014), 36. 2, 7
- [KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In *Proc. CVPR* (2015). 2
- [KLCF10] KIM V. G., LIPMAN Y., CHEN X., FUNKHOUSER T. A.: Möbius transformations for global intrinsic symmetry analysis. *Comput. Graph. Forum* 29, 5 (2010), 1689–1700. 1
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T. A.: Blended intrinsic maps. *TOG* 30, 4 (2011), 79. 1, 9
- [LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. In *Proc. SGP* (2008), pp. 1421–1430. 1

- [Mac49] MACNEAL R. H.: *The solution of partial differential equations by means of electrical networks*. PhD thesis, California Institute of Technology, 1949. 3
- [MH88] MURTHY D. V., HAFTKA R. T.: Derivatives of eigenvalues and eigenvectors of a general complex matrix. *Intl. J. Numer. Met. Eng.* 26, 2 (1988), 293–311. 4, 6, 14
- [MS89] MUMFORD D., SHAH J.: Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure and Applied Math.* 42, 5 (1989), 577–685. 7
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4 (July 2012), 30:1–30:11. 2, 3, 7, 8, 9, 10
- [PBB13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M.: Partial shape matching without point-wise correspondence. *Numer. Math. Theor. Meth. Appl.* 6 (2013), 223–244. 2, 7, 9
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental mathematics* 2, 1 (1993), 15–36. 3
- [RABT13] RODOLÀ E., ALBARELLI A., BERGAMASCO F., TORSELLO A.: A scale independent selection process for 3d object recognition in cluttered scenes. *International Journal of Computer Vision* 102, 1-3 (2013), 129–145. 11
- [RBA*12] RODOLÀ E., BRONSTEIN A., ALBARELLI A., BERGAMASCO F., TORSELLO A.: A game-theoretic approach to deformable shape matching. In *Proc. CVPR* (June 2012), pp. 182–189. 2
- [RTH*13] RODOLÀ E., TORSELLO A., HARADA T., KUNIOYOSHI Y., CREMERS D.: Elastic net constraints for shape matching. In *Proc. ICCV* (December 2013), pp. 1169–1176. 2, 9
- [SY14] SAHILIOĞLU Y., YEMEZ Y.: Partial 3-d correspondence from shape extremities. *Computer Graphics Forum* 33, 6 (2014), 63–76. 2, 9
- [TSDS10] TOMBARI F., SALTI S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *Proc. ECCV* (2010), pp. 356–369. 9
- [VC02] VESE L. A., CHAN T. F.: A multiphase level set framework for image segmentation using the Mumford and Shah model. *IJCV* 50, 3 (2002), 271–293. 7
- [VKTS*11] VAN KAICK O., TAGLIASACCHI A., SIDI O., ZHANG H., COHEN-OR D., WOLF L., HAMARNEH G.: Prior knowledge for part correspondence. *Computer Graphics Forum* 30, 2 (2011), 553–562. 2
- [VKZH13] VAN KAICK O., ZHANG H., HAMARNEH G.: Bilateral maps for partial matching. *Computer Graphics Forum* 32, 6 (2013), 189–200. 2, 11
- [VKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. 1
- [Wey11] WEYL H.: Über die asymptotische verteilung der eigenwerte. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse 1911* (1911), 110–117. 4
- [WSSC11] WINDHEUSER T., SCHLICKWEI U., SCHMIDT F. R., CREMERS D.: Large-scale integer linear programming for orientation preserving 3d shape matching. *Computer Graphics Forum* 30, 5 (2011), 1471–1480. 2

Appendix A - Discretization

We describe more in depth the discretization steps and give the implementation details of our algorithm (we skip trivial derivations for

the sake of compactness). Detailed gradients for each term are given in Appendix B.

Mesh parametrization. We denote by S a triangle mesh of n points, composed of triangles S_j for $j = 1, \dots, m$. In the following derivations, we will consider the classical triangle-based parametrization described by the charts $x_j: \mathbb{R}^2 \rightarrow \mathbb{R}^3$

$$x_j(\alpha, \beta) = x_{j,1} + \alpha(x_{j,2} - x_{j,1}) + \beta(x_{j,3} - x_{j,1}), \quad (18)$$

with $\alpha \in [0, 1]$ and $\beta \in [0, 1 - \alpha]$. With $x_{j,k} \in \mathbb{R}^3$ we denote the 3D coordinates of vertex $k \in \{1, 2, 3\}$ in triangle S_j .

Each triangle S_j is equipped with a discrete metric tensor with coefficients

$$g_j = \begin{pmatrix} E_j & F_j \\ F_j & G_j \end{pmatrix}, \quad (19)$$

where $E_j = \|x_{j,2} - x_{j,1}\|^2$, $F_j = \langle x_{j,2} - x_{j,1}, x_{j,3} - x_{j,1} \rangle$, and $G_j = \|x_{j,3} - x_{j,1}\|^2$. The volume element for the j -th triangle is then given by $\sqrt{\det g_j} = \sqrt{E_j G_j - F_j^2}$.

Integral of a scalar function. Scalar functions $f: S \rightarrow \mathbb{R}$ are assumed to behave linearly within each triangle. Hence, $f(x(\alpha, \beta))$ is a linear function of (α, β) and it is uniquely determined by its values at the vertices of the triangle. The integral of f over S_j is then simply given by:

$$\begin{aligned} & \int_0^1 \int_0^{1-\alpha} f(\alpha, \beta) \sqrt{\det g_j} d\beta d\alpha \\ &= \int \int f(0,0)(1-\alpha-\beta) + f(1,0)\alpha + f(0,1)\beta \sqrt{\det g_j} d\beta d\alpha \\ &= \frac{1}{6}(f(0,0) + f(1,0) + f(0,1)) \sqrt{E_j G_j - F_j^2} \\ &= \frac{1}{3}(f(0,0) + f(1,0) + f(0,1)) \text{area}(S_j), \end{aligned} \quad (20)$$

where $f(0,0) = f(x_{j,1})$, $f(1,0) = f(x_{j,2})$, and $f(0,1) = f(x_{j,3})$.

Gradient of a scalar function. For the intrinsic gradient of f we get the classical expression in local coordinates:

$$\nabla f = \begin{pmatrix} \frac{\partial x_j}{\partial \alpha} & \frac{\partial x_j}{\partial \beta} \end{pmatrix} \begin{pmatrix} E_j & F_j \\ F_j & G_j \end{pmatrix}^{-1} \begin{pmatrix} f_\alpha \\ f_\beta \end{pmatrix}, \quad (22)$$

where we write f_α to denote the partial derivative $\frac{\partial f}{\partial \alpha} = f_{j,2} - f_{j,1}$ and similarly for f_β . The norm of the intrinsic gradient over triangle S_j is then given by:

$$\begin{aligned} & \|\nabla f\| \\ &= \sqrt{\langle \nabla f, \nabla f \rangle} \\ &= \sqrt{\begin{pmatrix} f_\alpha & f_\beta \end{pmatrix} \begin{pmatrix} E_j & F_j \\ F_j & G_j \end{pmatrix}^{-1} \begin{pmatrix} f_\alpha \\ f_\beta \end{pmatrix}} \\ &= \sqrt{\begin{pmatrix} f_\alpha & f_\beta \end{pmatrix} \begin{pmatrix} G_j & -F_j \\ -F_j & E_j \end{pmatrix} \begin{pmatrix} f_\alpha \\ f_\beta \end{pmatrix} \frac{1}{\sqrt{\det g_j}}} \\ &= \sqrt{\frac{f_\alpha^2 G_j - 2f_\alpha f_\beta F_j + f_\beta^2 E_j}{\det g_j}}. \end{aligned} \quad (23)$$

Note that, since we take f to be linear, the gradient ∇f is constant within each triangle. We can then integrate ∇f over S_j as follows:

$$\begin{aligned} & \int_{S_j} \|\nabla f(x)\| \sqrt{\det g_j} d\alpha d\beta & (25) \\ &= \int_{S_j} \sqrt{\frac{f_\alpha^2 G_j - 2f_\alpha f_\beta F_j + f_\beta^2 E_j}{\det g_j}} \sqrt{\det g_j} d\alpha d\beta \\ &= \int_{S_j} \sqrt{f_\alpha^2 G_j - 2f_\alpha f_\beta F_j + f_\beta^2 E_j} d\alpha d\beta \\ &= \frac{1}{2} \sqrt{f_\alpha^2 G_j - 2f_\alpha f_\beta F_j + f_\beta^2 E_j}. & (26) \end{aligned}$$

In the following, we write \mathcal{N} and \mathcal{M} to denote the partial and full shape respectively. Further, let $\{\lambda_i^{\mathcal{N}}\}_{i=1,\dots,k}$ be the first k eigenvalues of the Laplacian on \mathcal{N} , and similarly for $\{\lambda_i^{\mathcal{M}}\}_{i=1,\dots,k}$. The functional map \mathbf{C} has size $k \times k$.

Mumford-Shah functional (μ_2 -term). Following Equations (21) and (26), we immediately obtain:

$$\begin{aligned} & \int_S \xi(v) \|\nabla v\| dx \\ &= \sum_{j=1}^m \int_{S_j} \xi(\alpha, \beta) \|\nabla v\| \sqrt{\det g_j} d\beta d\alpha \\ &= \sum_{j=1}^m \sqrt{v_\alpha^2 G_j - 2v_\alpha v_\beta F_j + v_\beta^2 E_j} \int_{S_j} \xi(\alpha, \beta) d\beta d\alpha \\ &\approx \frac{1}{6} \sum_{j=1}^m \sqrt{v_\alpha^2 G_j - 2v_\alpha v_\beta F_j + v_\beta^2 E_j} (\xi(0,0) + \xi(1,0) + \xi(0,1)), \end{aligned}$$

where $\xi(0,0) = \xi(v(x_{j,1}))$, $\xi(1,0) = \xi(v(x_{j,2}))$, and $\xi(0,1) = \xi(v(x_{j,3}))$.

Weight matrix (μ_3 -term). Recall from Section 3 and Figure 6 that an estimate for the rank of \mathbf{C} can be easily computed as

$$r = \max\{i \mid \lambda_i^{\mathcal{N}} < \max_j \lambda_j^{\mathcal{M}}\}. \quad (27)$$

We use this information in order to construct the weight matrix \mathbf{W} , whose diagonal slope directly depends on r .

To this end, we model \mathbf{W} as a regular $k \times k$ grid in \mathbb{R}^2 . The slanted diagonal of \mathbf{W} is a line segment $\delta(t) = \mathbf{p} + t \frac{\mathbf{n}}{\|\mathbf{n}\|}$ with $t \in \mathbb{R}$, where $\mathbf{p} = (1, 1)^\top$ is the matrix origin, and $\mathbf{n} = (1, r/k)^\top$ is the line direction with slope r/k . The high-frequency spread in \mathbf{C} is further accounted for by funnel-shaping \mathbf{W} along the slanted diagonal. We arrive at the following expression for \mathbf{W} :

$$w_{ij} = e^{-\sigma \sqrt{r^2 + j^2}} \frac{\mathbf{n}}{\|\mathbf{n}\|} \times ((i, j)^\top - \mathbf{p}), \quad (28)$$

where the second factor is the distance from the slanted diagonal δ , and $\sigma \in \mathbb{R}_+$ regulates the spread around δ . In our experiments we set $\sigma = 0.03$.

Orthogonality (μ_4, μ_5 -terms). For practical reasons, we incorporate the off-diagonal and diagonal terms within one term with the single coefficient $\mu_{4,5}$. In addition, we rewrite the off-diagonal penalty using the following equivalent expression:

$$\sum_{i \neq j} (\mathbf{C}^\top \mathbf{C})_{ij}^2 = \|\mathbf{C}^\top \mathbf{C}\|_F^2 - \sum_i (\mathbf{C}^\top \mathbf{C})_{ii}^2. \quad (29)$$

submitted to COMPUTER GRAPHICS Forum (12/2015).

Vector $\mathbf{d} \in \mathbb{R}^k$ is constructed by setting the first r elements (according to (27)) equal to 1, and the remaining $k - r$ elements equal to 0.

Appendix B - Gradients

We find local solutions to each optimization problem by the (nonlinear) conjugate gradient method. In this Section we give the detailed gradient derivations of all terms involved in the optimization.

In order to keep the derivations practical, we will model function v by its corresponding n -dimensional vector \mathbf{v} . Note that, depending on the optimization step, the gradients are computed with respect to either \mathbf{v} or \mathbf{C} .

Data term (w.r.t. v). Let q denote the number of corresponding functions between the two shapes. Matrices \mathbf{F} and \mathbf{G} contain the column-stacked functions defined over \mathcal{N} and \mathcal{M} and have size $n \times q$. The respective projections onto the corresponding functional spaces Φ and Ψ are stored in the $k \times q$ matrices \mathbf{A} and \mathbf{B} respectively. Let us write \mathbf{H}_{ij} to identify the elements of the matrix $\mathbf{CA} - \mathbf{B}(\eta(v))$, we then have

$$\begin{aligned} & \frac{\partial}{\partial v_p} \|\mathbf{CA} - \mathbf{B}(\eta(v))\|_{2,1} \\ &= \frac{\partial}{\partial v_p} \sum_{j=1}^q \left(\sum_{i=1}^n \mathbf{H}_{ij}^2 \right)^{\frac{1}{2}} \\ &= \sum_{j=1}^q \left(\sum_{i=1}^n \mathbf{H}_{ij}^2 \right)^{-\frac{1}{2}} \sum_{i=1}^n \mathbf{H}_{ij} \frac{\partial}{\partial v_p} \mathbf{H}_{ij}. \end{aligned} \quad (30)$$

Since $\mathbf{B}(\eta(v))_{ij} = \sum_k \Psi_{ik}^\top \eta(v_k) \mathbf{F}_{kj}$, we have:

$$\frac{\partial}{\partial v_p} \mathbf{H}_{ij} = \frac{\partial}{\partial v_p} [\mathbf{CA}_{ij} - \mathbf{B}_{ij}] = \Psi_{ip}^\top \mathbf{F}_{pj} \frac{\partial}{\partial v_p} \eta(v_p). \quad (31)$$

Finally:

$$\frac{\partial}{\partial v_p} \eta(v_p) = 1 - \tanh^2(2v_p - 1). \quad (32)$$

Area term (μ_1 -term w.r.t. v). The derivative of the discretized area term is:

$$\begin{aligned} & \frac{\partial}{\partial v_p} \left(\sum_{i=1}^n (S_{\mathcal{N}})_i - \sum_{i=1}^n (S_{\mathcal{M}})_i \eta(v_i) \right)^2 \\ &= 2 \left(\sum_{i=1}^n (S_{\mathcal{N}})_i - \sum_{i=1}^n (S_{\mathcal{M}})_i \eta(v_i) \right) S_M \frac{\partial}{\partial v_p} \eta(v_p) \end{aligned}$$

where $(S_{\mathcal{M}})_i$ and $(S_{\mathcal{N}})_i$ are the local area elements associated with the i -th vertex of meshes \mathcal{M} and \mathcal{N} respectively. For the derivative of $\eta(v_p)$ see equation (32).

Mumford-Shah functional (μ_2 -term w.r.t. v). Computing the gradient $\nabla_v \int_S \xi(v) \|\nabla v\| dx$ involves computing partial derivatives of $\xi(v)$ with respect to \mathbf{v} . These are simply given by:

$$\begin{aligned} \frac{\partial}{\partial v_k} \xi(v_k) &= \frac{\partial}{\partial v_k} e^{-\frac{\tanh(2v_k - 1)}{4\sigma^2}} \\ &= -\frac{1 - \tanh^2(2v_k - 1)}{2\sigma^2} e^{-\frac{\tanh(2v_k - 1)}{4\sigma^2}}. \end{aligned}$$

In the following derivations we set $D_j \equiv \sqrt{v_\alpha^2 G_j - 2v_\alpha v_\beta F_j + v_\beta^2 E_j}$, and $D_j = 0$ whenever $\nabla \mathbf{v} = \mathbf{0}$. The gradient of the Mumford-Shah functional is then composed of the partial derivatives:

$$\begin{aligned} & \frac{\partial}{\partial v_k} \int_S \xi(\mathbf{v}) \|\nabla \mathbf{v}\| dx \\ &= \sum_{j=1}^m \frac{\partial}{\partial v_k} \int_{S_j} \xi(\mathbf{v}) \|\nabla \mathbf{v}\| \\ &= \frac{1}{6} \sum_{j \in N(k)} \frac{\partial}{\partial v_k} D_j (\xi(v_k) + \xi(v_{j,2}) + \xi(v_{j,3})) \\ &= \frac{1}{6} \sum_{j \in N(k)} (\xi(v_k) + \xi(v_{j,2}) + \xi(v_{j,3})) \frac{\partial}{\partial v_k} D_j + D_j \frac{\partial}{\partial v_k} \xi(v_k) \\ &= \frac{1}{6} \sum_{j \in N(k)} (\xi(v_k) + \xi(v_{j,2}) + \xi(v_{j,3})) \frac{1}{2D_j} \frac{\partial}{\partial v_k} D_j^2 + D_j \frac{\partial}{\partial v_k} \xi(v_k) \\ &= \frac{1}{6} \sum_{j \in N(k)} (\xi(v_k) + \xi(v_{j,2}) + \xi(v_{j,3})) K_j + D_j \frac{\partial}{\partial v_k} \xi(v_k) \end{aligned}$$

where we write $K_j \equiv \frac{1}{D_j} ((v_k - v_{j,2})(G_j - F_j) + (v_k - v_{j,3})(E_j - F_j))$, and $j \in N(k)$ are the indices of the triangles containing the k -th vertex. Note that we slightly abuse notation by writing $v_k, v_{j,2}$, and $v_{j,3}$ to denote the three vertices of the j -th triangle, even though in general the ordering might be different depending on the triangle.

Data term (w.r.t. C). The derivative of the data term with respect to \mathbf{C} is similar to (30). The only difference is in the partial derivative:

$$\frac{\partial}{\partial C_{pq}} \mathbf{H}_{ij} = \frac{\partial}{\partial C_{pq}} \sum_{i=1}^k \mathbf{C}_{ik} \mathbf{A}_{kj} = \begin{cases} \mathbf{C}_{pq} \mathbf{A}_{qj} & \text{if } i = p \\ 0 & \text{otherwise.} \end{cases} \quad (33)$$

Weight matrix (μ_3 -term w.r.t. C). This is simply given by:

$$\frac{\partial}{\partial C_{pq}} \|\mathbf{C} \circ \mathbf{W}\|_{\mathbb{F}}^2 = \mathbf{C}_{pq} (\mathbf{W}_{pq})^2. \quad (34)$$

Orthogonality ($\mu_{4,5}$ -term w.r.t. C). The gradient of the last term can be finally obtained as:

$$\begin{aligned} & \frac{\partial}{\partial C_{pq}} \left[\|\mathbf{C}^T \mathbf{C}\|_{\mathbb{F}}^2 - \sum_i (\mathbf{C}^T \mathbf{C})_{ii}^2 + \sum_i ((\mathbf{C}^T \mathbf{C})_{ii} - d_i)^2 \right] \\ &= 4(\mathbf{C} \mathbf{C}^T \mathbf{C})_{pq} + 2 \sum_i \left[-\sum_k \mathbf{C}_{ki}^2 \frac{\partial}{\partial C_{pq}} \sum_k \mathbf{C}_{ki}^2 + (\sum_k \mathbf{C}_{ki}^2 - d_i) \frac{\partial}{\partial C_{pq}} \sum_k \mathbf{C}_{ki}^2 \right] \\ &= 4[(\mathbf{C} \mathbf{C}^T \mathbf{C})_{pq} - d_q \mathbf{C}_{pq}]. \end{aligned}$$

Appendix C - Perturbation Analysis

Theorem 1 Let $\mathbf{L}_{\mathcal{N}} + t\mathbf{P}_{\mathcal{N}} = \Phi(t)^T \mathbf{A}(t) \Phi(t)$, where $\mathbf{A}(t) = \text{diag}(\lambda_1(t), \dots, \lambda_n(t))$ is a diagonal matrix of eigenvalues, and $\Phi(t)$ are the corresponding eigenvectors. The derivative of the non-trivial eigenvalues is given by

$$\frac{d}{dt} \lambda_i = \sum_{v,w \in \partial \mathcal{N}} (\mathbf{P}_{\mathcal{N}})_{vw} \phi_{iv} \phi_{iw} = \Phi_i^T \mathbf{P}_{\mathcal{N}} \Phi_i. \quad (35)$$

Proof: Let $\mathbf{A}(t)$ be a symmetric real $n \times n$ matrix parametrized by $t \in T \subseteq \mathbb{R}$, with $\Phi(t)$ and $\mathbf{A}(t)$ being the eigenvector and eigenvalue matrices, i.e., for all $t \in T$ we have

$$\mathbf{A}(t) \Phi(t) = \Phi(t) \mathbf{A}(t) \quad (36)$$

and $\mathbf{A}(t)$ is diagonal and $\Phi(t)$ orthogonal.

Following [MH88], if all the eigenvalues are distinct, then we can compute the derivatives of the eigenvalues at $t = 0$ as

$$\lambda'_i = \Phi_i^T \mathbf{A}' \Phi_i \quad (37)$$

where \mathbf{A}' , the derivative of $\mathbf{A}(t)$, and the eigenvectors Φ_i are considered at $t = 0$. In fact, differentiating (35), we obtain

$$\mathbf{A}' \Phi + \mathbf{A} \Phi' = \Phi' \mathbf{A} + \Phi \mathbf{A}'. \quad (38)$$

Left-multiplying both sides by Φ^T , setting $\Phi' = \Phi \mathbf{B}$ for a matrix \mathbf{B} to be determined, and recalling that $\Phi^T \mathbf{A} \Phi = \mathbf{A}$, we have

$$\Phi^T \mathbf{A}' \Phi + \mathbf{A} \mathbf{B} = \mathbf{B} \mathbf{A} + \mathbf{A}', \quad (39)$$

from which

$$\text{diag}(\mathbf{A}') = \text{diag}(\Phi^T \mathbf{A}' \Phi) + \text{diag}(\mathbf{A} \mathbf{B} - \mathbf{B} \mathbf{A}) = \text{diag}(\Phi^T \mathbf{A}' \Phi). \quad (40)$$

Going back to our case, we take the simplifying assumptions that $\mathbf{L}_{\mathcal{M}}$ and $\mathbf{L}_{\mathcal{N}}$ do not have repeated eigenvalues. Let

$$\mathbf{L}_{\mathcal{N}} = \Phi^T \mathbf{A} \Phi \quad (41)$$

$$\mathbf{L}_{\overline{\mathcal{N}}} = \overline{\Phi}^T \overline{\mathbf{A}} \overline{\Phi} \quad (42)$$

be the spectral decompositions of $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\overline{\mathcal{N}}}$ respectively. According to the previous result, we can write the derivative of λ_i eigenvalue of $\mathbf{L}_{\mathcal{N}}$ and, thus, of $\mathbf{L}(0)$, as:

$$\lambda'_i = \Phi_i^T \mathbf{P}_{\mathcal{N}} \Phi_i = \sum_{v,w \in \partial \mathcal{N}} (\mathbf{P}_{\mathcal{N}})_{vw} \phi_{iv} \phi_{iw}. \quad (43)$$

Theorem 2 Assume that $\mathbf{L}_{\mathcal{N}}$ has distinct eigenvalues ($\lambda_i \neq \lambda_j$ for $i \neq j$), and furthermore, the non-zero eigenvalues are all distinct from the eigenvalues of $\mathbf{L}_{\overline{\mathcal{N}}}$ ($\lambda_i \neq \bar{\lambda}_j$ for all i, j). Let $\mathbf{L}_{\mathcal{N}} + t\mathbf{P}_{\mathcal{N}} = \Phi(t)^T \mathbf{A}(t) \Phi(t)$, where $\mathbf{A}(t) = \text{diag}(\lambda_1(t), \dots, \lambda_n(t))$ is a diagonal matrix of eigenvalues, and $\Phi(t)$ are the corresponding eigenvectors. Then, the derivative of the non-constant eigenvector is given by

$$\frac{d}{dt} \Phi_i = \sum_{j=1}^n \frac{\Phi_i^T \mathbf{P}_{\mathcal{N}} \Phi_j}{\lambda_i - \lambda_j} \Phi_j + \sum_{j=1}^{\bar{n}} \frac{\Phi_i^T \mathbf{P} \overline{\Phi}_j}{\lambda_i - \bar{\lambda}_j} \overline{\Phi}_j. \quad (44)$$

Proof: Under the same assumptions as for the previous theorem, from (39) we have

$$(\Phi^T \mathbf{A}' \Phi)_{ij} + (\mathbf{A} \mathbf{B})_{ij} = (\mathbf{B} \mathbf{A})_{ij} + (\mathbf{A}')_{ij} \quad (45)$$

$$\Phi_i^T \mathbf{A}' \Phi_j + \lambda_i b_{ij} = b_{ij} \lambda_j + 0, \quad (46)$$

from which

$$b_{ij} = \frac{\Phi_i^T \mathbf{A}' \Phi_j}{\lambda_j - \lambda_i}. \quad (47)$$

due to the orthogonality of $\Phi(t)$ we have that \mathbf{B} is skew-symmetric, and thus, $b_{ii} = 0$. From the relation $\Phi' = \Phi \mathbf{B}$ we obtain

$$\frac{d}{dt} \Phi_i = \sum_{j \neq i} b_{ji} \Phi_j = \sum_{j \neq i} \frac{\Phi_j^T \mathbf{A}' \Phi_i}{\lambda_i - \lambda_j} \Phi_j. \quad (48)$$

Going back to our case, recall that the set of eigenvalues of $\mathbf{L}(0)$ is the union of the eigenvalues of $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\overline{\mathcal{N}}}$ and the corresponding eigenvectors are obtained from those of $\mathbf{L}_{\mathcal{N}}$ and $\mathbf{L}_{\overline{\mathcal{N}}}$ by padding with zeros on the missing parts. Denoting with λ_i and Φ_i the eigenvalues and corresponding eigenvectors of $\mathbf{L}_{\mathcal{N}}$, and $\bar{\lambda}_j$ and $\overline{\Phi}_j$ the

eigenvalues and corresponding eigenvectors of $\mathbf{L}_{\overline{\mathcal{N}}}$, we have

$$\frac{d}{dt} \phi_i = \sum_{\substack{j=1 \\ j \neq i}}^n \frac{\phi_i^\top \mathbf{P}_{\mathcal{N}} \phi_j}{\lambda_i - \lambda_j} \phi_j + \sum_{j=1}^{\bar{n}} \frac{\phi_i^\top \mathbf{P} \bar{\phi}_j}{\lambda_i - \bar{\lambda}_j} \bar{\phi}_j. \quad (49)$$

Boundary interaction strength. We can measure the variation of the eigenbasis as a function of the boundary \mathcal{B} splitting \mathcal{M} into \mathcal{N} and $\overline{\mathcal{N}}$ as

$$\partial \Phi(\mathcal{B}) = \sum_{i=1}^n \|\phi_i'\|_{\mathcal{N}}^2 = \sum_{i=1}^n \left(\sum_{\substack{j=1 \\ j \neq i}}^n \frac{\phi_i^\top \mathbf{D}_{\mathcal{N}} \phi_j}{\lambda_i - \lambda_j} \right)^2. \quad (50)$$

Let us now consider the function:

$$f(v) = \sum_{\substack{i,j=1 \\ j \neq i}}^n \left(\frac{\phi_{iv} \phi_{jv}}{\lambda_i - \lambda_j} \right)^2. \quad (51)$$

Assuming $\mathbf{D}_{\mathcal{N}}$ diagonal and with constant diagonal elements k , we have

$$k \int_{\mathcal{B}} f(v) dv \geq \partial \Phi(\mathcal{B}), \quad (52)$$

in fact:

$$\begin{aligned} \partial \Phi(\mathcal{B}) &\approx k \sum_{i=1}^n \left(\sum_{\substack{j=1 \\ j \neq i}}^n \frac{\sum_{v \in \partial \mathcal{M}} \phi_{iv} \phi_{jv}}{\lambda_i - \lambda_j} \right)^2 \\ &\leq k \sum_{v \in \partial \mathcal{M}} \sum_{\substack{i,j=1 \\ j \neq i}}^n \left(\frac{\phi_{iv} \phi_{jv}}{\lambda_i - \lambda_j} \right)^2 = k \sum_{v \in \partial \mathcal{M}} f(v). \end{aligned} \quad (53)$$